

(19) World Intellectual Property Organization
International Bureau



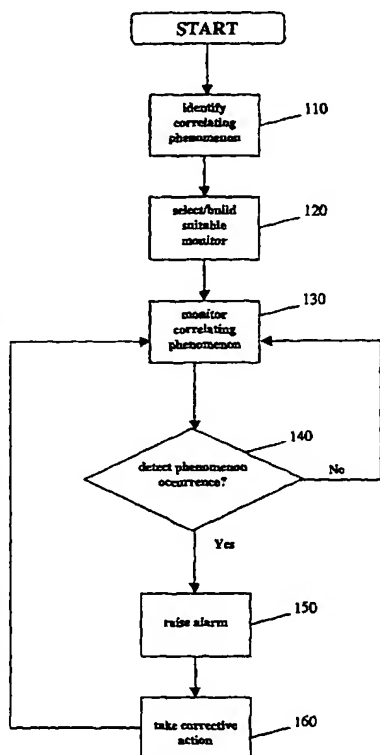
(43) International Publication Date
24 January 2002 (24.01.2002)

PCT

(10) International Publication Number
WO 02/06972 A1

- (51) International Patent Classification⁷: **G06F 15/173** (72) Inventor; and
(75) Inventor/Applicant (for US only): **LEWIS, Lundy, M.**
(21) International Application Number: **PCT/US01/22108** [US/US]; 480 Greenville Road, Mason, NH 03048 (US).
(74) Agent: **HENDERSON, Norma, E.**; 13 Jefferson Drive,
Londonderry, NH 03053 (US).
(25) Filing Language: English (81) Designated States (national): AU, CA, US.
(26) Publication Language: English (84) Designated States (regional): European patent (AT, BE,
CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC,
NL, PT, SE, TR).
(30) Priority Data: 60/217,972 13 July 2000 (13.07.2000) US
Published:
— with international search report
(71) Applicant (for all designated States except US):
APRISMA MANAGEMENT TECHNOLOGIES, INC. [US/US]; 121 Technology Drive, Durham, NH
03824 (US).
For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: **METHOD AND APPARATUS FOR MONITORING AND MAINTAINING USER-PERCEIVED QUALITY OF SERVICE IN A COMMUNICATIONS NETWORK**



(57) Abstract: A method and apparatus for managing data, voice, application, and video services allows anticipation of poor quality of service from a remote management station, in order to allow correction of the cause before the end user perceives service quality degradation. Specific system phenomena are identified (110) that coincide with user-perceived service degradation in a particular network. The network is then monitored for the occurrence of those phenomena (120). Incipient or existing user-perceived quality of service degradation is inferred from the occurrence of one or more of those phenomena (130, 140) and action is taken to avoid and/or correct the degraded service quality condition (150, 160). In a preferred embodiment, as many of the steps as possible are performed automatically by a network management system. In one embodiment, a close correlation is assumed between application data buffer over-extension and poor quality of service from a user's point of view. In this embodiment, a monitor (520) is placed on the application data buffer that raises an alarm for a network management system (540) whenever the buffer is close to over-extension (530) or an algorithm identifies a trend towards over-extension.

WO 02/06972 A1

**Method and Apparatus for Monitoring and Maintaining User-Perceived
Quality of Service in a Communications Network**

5 Field of the invention

The invention relates to management of communications networks and, in particular, to anticipation and avoidance of user-perceived service quality degradation.

10 Background

Typically, managers of communications networks exist far away from network services as experienced by end users. It is therefore difficult for these remotely situated managers to know when the quality of service, as perceived by the users, is unacceptable. Currently, a user usually must call a help desk if
15 the quality of service becomes unacceptable. Consequently, user work is interrupted and the achievement of the purpose of the service may be degraded or even halted for lengthy periods of time.

Objects of the invention

20 The object of the present invention is to provide a method and apparatus by which to allow, from a remote management station, anticipation of the onset of poor quality of video, voice, application, or data services in order allow correction of the cause or causes before the end user perceives a degradation in service quality.

25

Summary

In the present invention, specific system phenomena are identified that coincide with user-perceived Quality of Service (QoS) degradation in a particular network and associated systems. Once one or more specific correlating phenomena are identified, one or more monitors for use in detecting the occurrence of the phenomena are selected and/or built. Each monitor is then installed at an appropriate place in the network or in a system application, in order to allow detection of any occurrences of the correlating phenomenon. The network is then monitored for the occurrence of those phenomena and incipient or existing user-perceived QoS degradation is inferred from an occurrence. When an occurrence is detected, an alarm is raised for the network manager's attention. Action can then be taken to avoid and/or correct the degraded service quality condition. In the preferred embodiment of the invention, as many of the steps as possible are performed automatically, preferably by a network management system.

One embodiment of the invention utilizes the close correlation between application data buffer over-extension and user-perceived poor quality of service. In this embodiment, the fullness of the application data buffer is monitored and an alarm is raised in a network management system whenever the buffer is close to over-extension or an algorithm identifies a trend towards over-extension.

Brief Description of the Drawings

Fig. 1 illustrates the operation of an embodiment of the method for monitoring and maintaining Quality of Service of the present invention;

Fig. 2 is a block diagram of an embodiment of an apparatus for monitoring and maintaining Quality of Service according to the present invention;

Fig. 3 illustrates the temporal sequence of the application data buffer re-flushing phenomenon;

Fig. 4 is a diagram of an illustrative network with which an embodiment of the present invention may be used; and

Fig. 5 illustrates the operation of an embodiment of the method of the present invention utilizing an application data buffer re-flushing monitor.

5

Detailed Description

With today's complex communications services involving video, data, and voice, there is evidence that an end user's perception of degraded quality of service (QoS) frequently coincides with specific, detectable, system phenomena. Such phenomena can include such things as CPU overload, near depletion of internal or external data stores, slow screen refreshing, and data buffer re-flushing. While it is not always clear whether these phenomena, either alone or in concert, are the cause of degraded service quality, are symptoms of the problem, or simply happen to generally coincide with degraded service quality conditions, their presence can still be utilized by network managers to anticipate and then avoid lengthy periods of degraded service quality.

The present invention involves identifying specific system phenomena that are related to user-perceived QoS degradation in a particular communications network and associated systems, monitoring that network and associated systems for the occurrence of those phenomena, inferring incipient or existing user-perceived QoS degradation from the occurrence of one or more of those phenomena, and taking action to avoid and/or correct a degraded service quality condition. In the preferred embodiment of the invention, as many of the steps as possible are performed automatically, e.g., once specific service quality degradation-related phenomena are identified, the network and associated systems are automatically monitored for their occurrence via a network management system and, when occurrences are detected, the corrective actions are automatically initiated by a network management system or other management apparatus, such as an element management system or a management agent.

The advantages of the invention over what has been done previously include providing the ability to infer the end user's perceived quality of service from remote management stations without user action, as well as to anticipate and prevent degradation of user-perceived quality of service. Collecting
5 variables constituting the state of a network and associated systems over time allows the use of machine learning algorithms to discover subtle causes of poor quality; such discoveries can then be used to increase the efficiency of the network.

An operational flowchart of the invention is shown in Fig. 1. In Fig. 1,
10 specific network phenomena that correlate with periods of degraded user-perceived Quality of Service are identified 110. In the preferred embodiment, a network management system (NMS) is used to gather data on system parameters and on occurrences of user-perceived QoS. The identification of specific system phenomena that coincide with periods of user-perceived QoS
15 degradation may be made through any number of methods known in the art including, but not limited to, statistical correlation, data mining algorithms, machine learning algorithms, reversing engineering application code or designs, and empirical observation. In the preferred embodiment, an NMS provides the correlation function, preferably being the same NMS used to
20 gather the data on which the correlation is performed.

Once one or more specific system phenomena related to user-perceived degraded QoS are identified 110, one or more monitors for use in detecting the occurrence of one or more of the phenomena are selected and/or built 120. Such monitors may include network management systems, management
25 agents, element management systems or any of the many other monitoring systems and devices known in the art. Monitoring of variables is accomplished through polling or traps, where the variables monitored may include buffer overflows, CPU overload, capacity of internal and external data stores, inferences from a collection number of such variables, or any of the
30 many other types of monitorable parameters known in the art. Each selected monitor is then installed in an appropriate place on the network or system

applications, in order to allow detection of any occurrences of the correlating phenomenon or phenomena 130. It is to be understood that the appropriate venue for the monitor depends on exactly what system phenomenon is being monitored and how.

5 When an occurrence of one of the correlating phenomena is detected 140, an alarm is raised 150 for the network manager's attention. In the preferred embodiment, each monitor and/or variable is automatically monitored by a network management system in order to facilitate the raising of alarms and taking of corrective actions. However, monitoring can of course be
10 handled by any of the many methods known in the art. In the preferred embodiment, the alarm is also raised automatically, again preferably by a network management system. Alarms may be raised by any of the many methods known in the art including, but not limited to, sending the alarm to a pager, to a telephone, to a network management system, to an element
15 management system, or to any other compatible system. Once the alarm is raised 150, corrective action is taken 160. In the preferred embodiment, this is also done automatically by a network management system, but it may alternatively be accomplished manually or by any of the other methods known in the art.

20 A block diagram of a system implementing the invention is shown in Fig. 2. As shown in Fig. 2, measurable network parameters and events 210 are compared to occurrences of user-perceived service degradation 220 by use of some form of correlation method 230. Any of the many correlation methods known in the art are suitable. In a preferred embodiment, the network
25 parameters 210 and occurrences of degraded user-perceived QoS 220 are collected by a network management system that is then used to perform correlation 230.

 Once a correlated phenomenon is identified through use of correlation method 230, occurrences of the phenomenon are monitored with phenomenon
30 monitor 240. In the preferred embodiment, this is also handled within a network manager. When an occurrence of the correlated phenomenon is

observed by the monitor 240, an alarm is raised by alarm-raising mechanism 250 and corrective action is taken by correction apparatus 260. In the preferred embodiment, alarm-raising mechanism 250 and correction apparatus 260 are again part of a network management system.

5 An example embodiment of the invention takes advantage of a system application phenomenon that has been observed to be related to user-perceived QoS degradation and is known as "data buffer re-flushing." In some systems, if a data buffer becomes over-extended (i.e. filled beyond its designated size), the buffer is flushed and begins to fill in again. This phenomenon has been
10 observed to coincide with flicks, specks, and irritating delays in the service (be it a voice, video, application, or data service).

 Fig. 3 illustrates the temporal sequence involved in application data buffer re-flushing. In Fig. 3, a data buffer is represented by an open-ended rectangle 302. The shaded area 304 represents the portion of the buffer that is
15 filled with data. Initially, the data buffer is only partially filled 310. As use of the application continues, the buffer becomes overfull 320. The buffer is then flushed 330 and begins to refill 340.

 In an embodiment of the present invention that makes use of the apparent correlation between data buffer re-flushing and service quality
20 degradation, a monitor is placed on the application data buffer in order that an alarm may be raised in a network management platform when (or just before) the data buffer is over-extended or an algorithm identifies a trend towards over-extension. In the preferred embodiment, the monitored variable would be a MIB variable (e.g. SNMP, CMIP, CORBA), but it could also be a variable
25 provided by a proprietary protocol or any other monitoring protocol. The monitor value at which an alarm is sent to the management system may be determined by a simple threshold function, a trending function, a fuzzy logic function, or any other appropriate function known in the art. The alarm may then be sent to a pager, a telephone, a network management system, an
30 element management system, or to any other compatible system.

An example of a networking scenario in which this particular embodiment is may be applied is the distance learning application. One distance learning application that has recently been studied for application of the present invention is part of the North Carolina Network Initiative (NCNI).

5 The NCNI participates in the "Internet2 Project", which refers to joint work among universities, industry, and federal agencies towards advancing Internet applications into spaces such as tele-medicine, remote laboratory work, and distance education.

At the core of the Internet2 design is a new technology referred to as a
10 Gigabit Point of Presence (GigaPoP). Given advances in fiber optic technology, the Internet backbone has become a more or less limitless, reliable medium for moving large volumes of traffic from one geographical area to another. A GigaPoP is the point of interconnection and service delivery between the institutional members of the Internet2 project and one or more
15 Internet service providers. GigaPoPs are essentially the on/off ramps between the Internet backbone and commercial businesses, university campuses, and government agencies. Thus, the GigaPoP is an intermediary network that regulates traffic between the Internet backbone and those other networks

The rationale for GigaPoP development is: Important as a very high-
20 performance backbone is to the next generation of Internet applications, it is no less important that the points at which people connect to the backbone, the so-called Points of Presence (PoPs), provide an equivalent level of performance. The Quality of Service of an Internet application, from the desktop, across the Internet, and back again, is only as good as the weakest
25 link in the application provision process. The requirement, then, is to build and manage a network that can serve as a PoP for handling the multi-gigabit traffic to be delivered by the next-generation Internet. The GigaPoP is a central distribution point where large amounts of digital traffic are moved between various end points and the main line. Since there will be diverse kinds of
30 applications that are downstream from the GigaPoP, each with special

bandwidth and priority requirements, it is important that the GigaPoP be able to regulate and prioritize traffic accordingly.

The Internet2 design calls for GigaPoPs that support several crucial features. Each GigaPoP must have high capacity (at least 622 Mb/s) and high
5 reliability and availability. It must use the Internet Protocol (IP) as a bearer service, and must be able to support emerging protocols and applications. It must be capable of serving simultaneously as a workaday environment and as a research test bed. It must allow for traffic measurement and data gathering. Lastly, it must permit migration to differentiated services and application-
10 aware networking.

NCNI built an intermediate GigaPoP network between the Internet2 backbone and the research community, with the goal of resolving bottlenecks in the community Internet typically caused by high traffic demands of distributed applications. The North Carolina GigaPoP is considered one of
15 several frontrunners in terms of research and development. Advanced applications such as distance education and remote laboratory work impose special requirements for managing the NC GigaPoP. The goals of the NC GigaPOP are (i) to keep local traffic local, (ii) to provide optimized access to research and education applications that depend upon the Internet and, most
20 importantly, (iii) to insure an acceptable quality of service for all local and Internet-driven applications, such as the distance learning application.

Figure 4 shows the overall topology of the NC GigaPoP. A GigaPoP is much like any other network, consisting of a collection of nodes. A node is a geographic location where various GigaPoP devices reside. As shown in Fig.
25 4, there are five primary nodes: North Carolina State University (NCSU) Centennial campus 402, NCSU Raleigh campus 404, Duke University 406, University of North Carolina at Chapel Hill 408, and MCNC 410. These primary nodes perform core routing and operational functions for the GigaPOP. They also serve as connection points (on-ramps) to national
30 networks, including the vBNS 420 and the Abilene network 422. Equipment at primary node sites includes optical-fiber terminating equipment 460, virtual-

circuit SONET switches 462, data routers 464, and network monitoring devices. Primary nodes are connected by OC48 SONET optical links 460. Secondary nodes reside at NCNI industry sites located in the RTP area, including Cisco Systems 430, IBM 432, and Nortel Networks 434. These
5 secondary nodes are connected to the primary nodes via tributary optical-fiber links 450. The NC GigaPoP is an intermediary network— the campus networks at NC State, Duke, and UNC at Chapel Hill are outside the scope of the GigaPoP but are connected to it

The North Carolina Research and Education Network (NCREN)
10 distance learning application was selected for experiments. Thus, the experiments focused specifically on those GigaPoP devices that support distance learning in North Carolina. In particular, the NC GigaPoP includes Litton Corporation CAMVision-2 Codec (CV2) video applications running on instructors' and students' NT workstations. CAMVision management is
15 therefore required in order to achieve (i) end-to-end management of the distance learning service and (ii) stronger event correlation and fault isolation over the complete set of elements that supports the distance learning service. It is equally clear that the quality and proactivity of CAMVision management has a great effect on the quality of service perceived by users of the distance
20 learning facility.

For this study, mapping of the network was limited to those Universities that were participating in the distance learning trials. The core router elements were added to the network management system being used, Aprisma Management Technologies' Spectrum®, using the "model by IP"
25 method. Spectrum retrieved MIB information from the routers, collected interface identifications and IP addresses, and discovered the logical and physical connections between the routers.

However, the end-to-end management of NC's distance learning application required some additional customization. While Spectrum has
30 management modules for other physical and logical NCREN objects, it was found to be necessary to develop a model in Spectrum that represents CV2s

and communicates with them via SNMP. It was determined that Litton's CAMVision (CV2) product runs on an NT workstation as an NT application; the Litton CV2 SNMP MIB piggybacks on the standard NT MIB. CV2, then, is basically an application that runs on an NT box. In Spectrum there is a
5 standard NT management module that provides the means to import additional NT application MIBs. This may be accomplished with the Spectrum Level-1 Toolkit, which means that it can be done on site, with no additional programming. Once the CV2/NT module is in place, each CV2 may be modeled for purposes of monitoring and control.

10 A correlation was performed between unexplained anomalies in the distance learning service and the state of the NC GigaPoP as a whole, and an investigation was undertaken to determine whether such knowledge could be used to answer the question: Can the network accommodate a particular new CAMVision video session and still meet the ultimate goal of proactive
15 management of user-perceived QoS? An investigation was also undertaken to determine the extent to which the information that could be acquired from CV2/NT was useful for management purposes. For example, were there CV2 MIB variables whose values indicate poor video quality? If so, then that would provide the means for an engineer to receive an alarm or page whenever poor
20 video quality occurs or is about to occur. Further, historical data were analyzed to infer conditions that typically coincide with poor video performance, including any recommendations for correction or possibly automated correction of the poor performance condition.

It was discovered that the only feedback regarding video quality that is
25 available at the application layer is the CV2 restart mechanism. That is, when the data buffers in a CV2 application are well beyond the full mark, the method of recovery is to flush the buffer and restart the sending stream. There was no trap in the CV2 MIB that allowed detection of an imminent restart condition. The insertion of a restart trap into the CV2 MIB is therefore useful
30 for management purposes, allowing full implementation of the present invention.

The hypothesis used for this embodiment of the present invention was, therefore, that the restart variable is the index into poor video quality; i.e. it correlates with poor quality of service from the user's point of view. User-perceived quality of service can therefore be inferred from a simple application
5 MIB variable — the buffer restart variable. The restart variable in fact proved to be a good index into quality of service from the user's point of view.

Enhancements were made to Spectrum in order to allow it to predict when an anomaly was about to happen and to then take action to prevent it. The Spectrum Alarm Manager provides various useful functions for the
10 implementation of the present invention, including automated popup when an alarm occurs and the ability to capture notes, probable causes, and other related data when acknowledging an alarm. Actually setting alarm thresholds and threshold formulas is largely straightforward using methods commonly known in the art. The more difficult task is deciding at what level to set
15 particular thresholds and developing advanced threshold formulas in the first place. This may be accomplished using any of a number of methods known in the art, including, but not limited to empirical experimentation, reverse engineering of application code or design, machine learning and statistical algorithms, and datamining. Spectrum was therefore set to raise an alarm at a
20 prespecified threshold-- for example, when the buffer restart variable is reset twice in less than a minute.

The operation of this embodiment of the present invention, utilizing a buffer re-flush restart monitor, is depicted in Fig. 5. In Fig. 5, a buffer re-flush monitor is built 510 by creating a variable in the CV2 MIB that reflects the
25 fullness of the video data buffers. This variable is monitored 520, and when it exceeds a prespecified threshold 530, an alarm is sent 540 to a remote network manager for corrective action.

The ultimate goal of the application of the example embodiment of the present invention to the distance learning service was to be able to predict poor
30 video performance and correct it before it occurs. Utilizing the present invention as part of end-to-end management of the distance learning service,

therefore, brings the distance learning application closer to the ultimate goal of proactive management of user-perceived Quality of Service. What has been described, however, is merely illustrative of the application of the principles of the present invention. Other arrangements, methods, modifications and
5 substitutions by one of ordinary skill in the art are also considered to be within the scope of the present invention, which is not to be limited except by the claims that follow.

CLAIMS

What is claimed is:

- 1 1. A method for remote detection of degraded Quality of Service in a
2 communications network and associated systems comprising, in combination, the
3 steps of:
4 identifying one or more system phenomena that coincide with the onset of
5 quality of service degradation;
6 monitoring said network for an occurrence of one of said phenomena; and
7 raising an alarm if an occurrence is detected.
2. The method of claim 1, further including the step of taking corrective
action to avoid said quality of service degradation.
3. The method of claim 1, wherein one or more of said steps of identifying,
monitoring, and raising are performed via a network management system.
4. The method of claim 2, wherein one or more of said steps of identifying,
monitoring, raising, and taking corrective action are performed via a network
management system.
- 1 5. An apparatus for remote detection of degraded Quality of Service in a
2 communications network and associated systems comprising, in combination:
3 means for identifying one or more system phenomena that coincide with the
4 onset of quality of service degradation;
5 at least one correlated phenomenon monitor for monitoring said network for
6 an occurrence of one or more of said phenomena; and
7 alarm-raising mechanism.
6. The apparatus of claim 5, further including means for taking corrective
action to avoid said quality of service degradation.

7. The apparatus of claim 5, wherein one or more of said means for identifying, said phenomenon monitor, and said alarm-raising mechanism are part of a network management system.

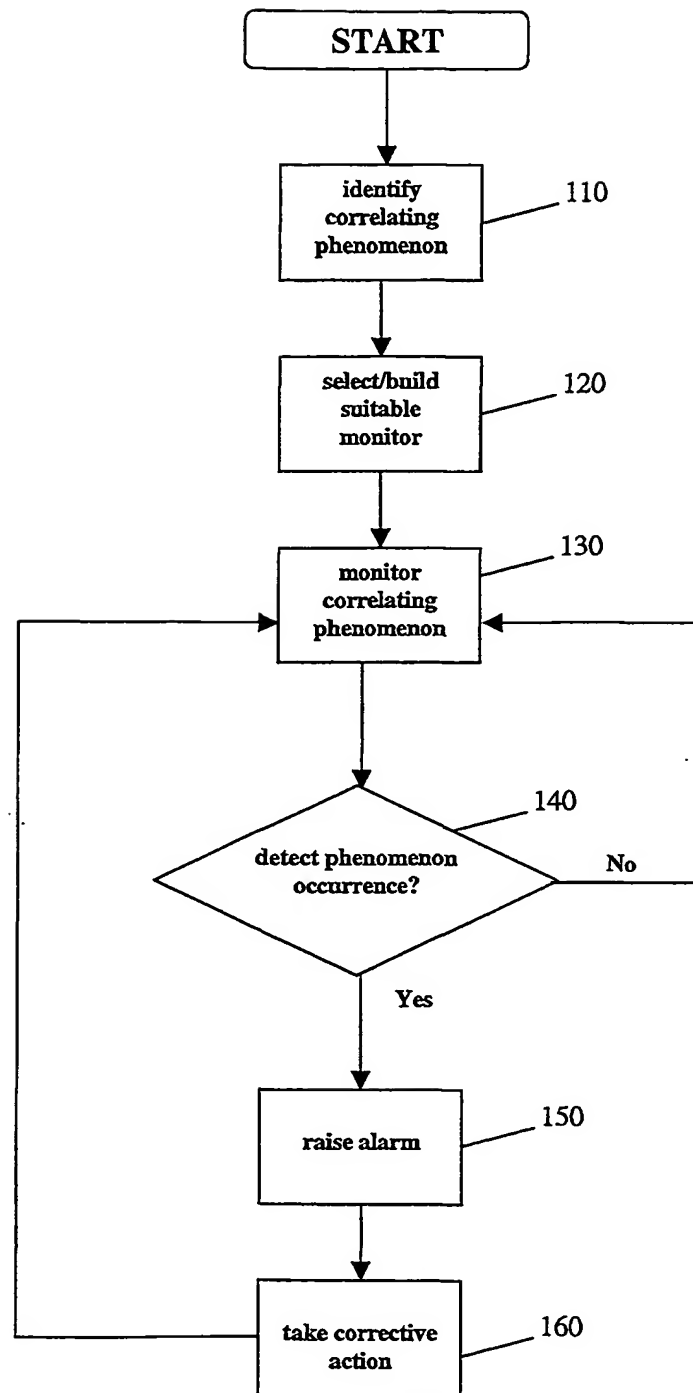
8. The apparatus of claim 6, wherein one or more of said means for identifying, said phenomenon monitor, said alarm-raising mechanism, and said means for taking corrective action are part of a network management system.

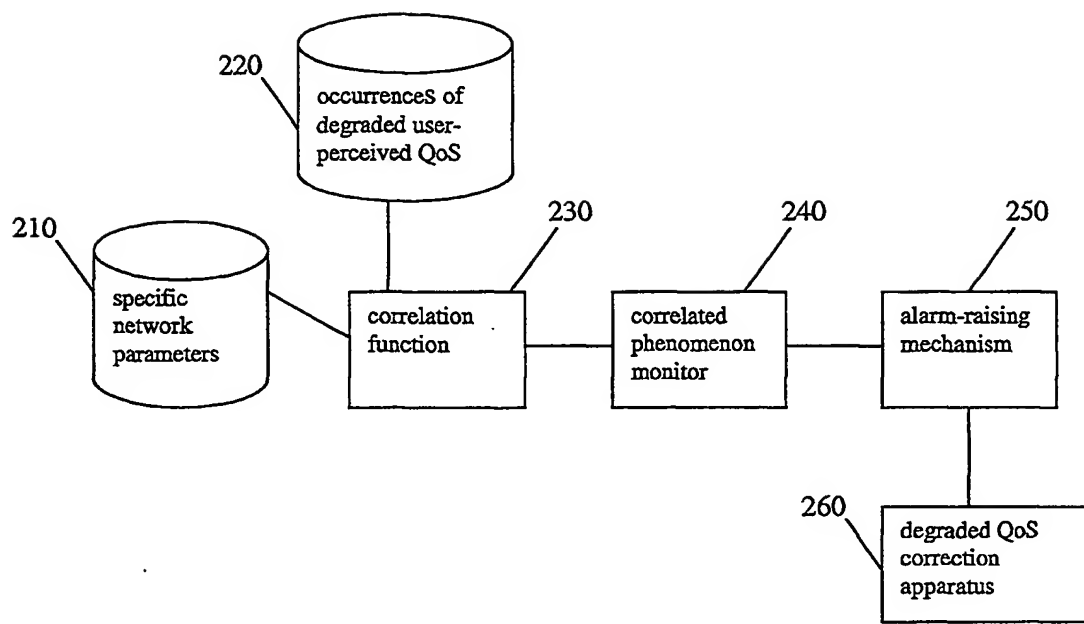
1 9. A method for remote detection of degraded Quality of Service in a
2 communications network and associated systems comprising, in combination, the
3 steps of:
4 establishing an application data buffer re-flush monitor;
5 monitoring said network for a threshold occurrence of application data buffer
6 re-flush; and
7 raising an alarm if a threshold occurrence is detected.


10. The method of claim 9, further including the step of taking corrective action to avoid said quality of service degradation.

11. The method of claim 9, wherein one or more of said steps of establishing, monitoring, and raising are performed via a network management system.

12. The method of claim 10, wherein one or more of said steps of establishing, monitoring, raising, and taking corrective action are performed via a network management system.

**Fig. 1**

**Fig. 2**

**Fig. 3**

North Carolina Networking Initiative




Fig. 4

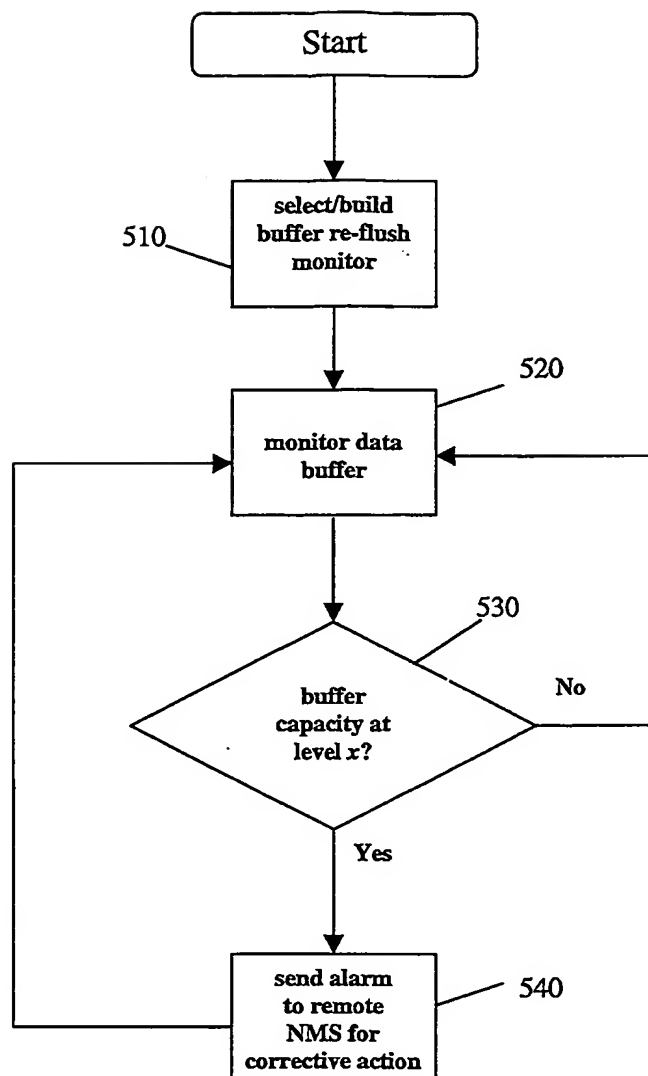


Fig. 5

INTERNATIONAL SEARCH REPORT

Int. application No.
PCT/US01/22108

A. CLASSIFICATION OF SUBJECT MATTER IPC(7) : G06F 15/173 US CL : 709/224, 226 According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED Minimum documentation searched (classification system followed by classification symbols) U.S. : 709/224, 226 370/231, 232; 714/47, 48 Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) IEEE ONLINE search terms: QOS, preventing, predicting, threshold		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5,898,668 A (SHAFFER) 27 April 1999, figures 2-4	1-12
A	5,703,870 A (MURASE) 30 December 1997, figure 2	1 and 7
X	US 6,012,152 A (DOUIK et al.) 04 January 2000, col. 13 and 33	1-8
X	US 5,130,978 A (MOBASSER) 14 July 1992, col. 6	1 and 7
X	LO et al. QOS Provisioning in Handoff Algorithms for Wireless LAN IEEE 1998, pages 9-16, especially pages 10-12	1-8
<input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/> See patent family annex.		
* Special categories of cited documents: "A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier document published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "&" document member of the same patent family	
Date of the actual completion of the international search 26 AUGUST 2001		Date of mailing of the international search report 28 SEP 2001
Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. (703) 305-3230		Authorized officer LARRY DONAGHUE <i>James R. Matthews</i> Telephone No. (703) 305-9675

INTERNATIONAL SEARCH REPORT

International application No.
PC1/US01/22108

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	MYERS Predicting and Measuring Quality of service for Mobile Multimedia, March 2000, pages 451-455	1-12